

kills, but you can pull in at come with Excel. For ad current prices for open a web query, click

Web into an Excel work-
b page, you can copy the
it into a refreshable web

Web button in the Get

paste the URL of the
et. Click the Go button

6. Click the Import button at the bottom of the New Web Query window to open the Import Data dialog box. The cell that you selected in step 1 is highlighted. Change this location, if necessary, and click the Properties button if you want to adjust the refresh times for the live data.
7. Click OK to add the web data to your worksheet. From the New Web Query window, if you had clicked the Save Query button, you could save the query settings as a separate file that you can open in another worksheet. After testing to ensure that the web data appears as expected, you can repeat these steps and save the query.
8. Repeat steps 1–7 to add data from another web page to your worksheet.

TIP FROM

EQ: Woody

If this process seems too cumbersome, you can speed things up by using the Clipboard. From an Internet Explorer window, select the data that you want to add to your worksheet, right-click, and choose Copy. Paste the data into your worksheet and click the Paste Options Smart Tag. Click the Create Refreshable Web Query menu option to convert the pasted data to a live link. This option is available only if you have selected a full table.

ANALYZING INFORMATION WITH PIVOTTABLES

PivotTables (as well as PivotCharts which you learn about in Chapter 22, “Creating and Editing Charts”) are powerful tools for automatically summarizing and analyzing data without ever having to add a formula or function. As the name implies, you start with a table, snap the rows and columns into position on a grid, and end up with a sorted, grouped, summarized, totaled, and subtotaled report. PivotTable reports are best for cross-tabulating tables—the more categories, the better. You can reduce a table of thousands of items to a single line, showing totals by category or quarter. Or you can create complex, multilevel groupings that show total sales by employee, grouped by product category and by quarter. You can hide or show detail for each group with a quick double-click. You can change the view or grouping in literally seconds, just by dragging items on or off the sheet and moving them between row, column, and page fields.

Start with a table that contains multiple fields, and then use Excel’s PivotTable button to set up a blank PivotTable page with just a few clicks. Instead of sorting your table and entering formulas and functions, you drag fields around on the PivotTable page to create a new view of your table—Excel groups the data and adds summary formulas automatically.

Unlike subtotals and outlines, which modify the structure of your table to display summaries, PivotTables create new, independent elements in your workbook. When you add or edit data in a table, the changes show up in your PivotTables (and PivotCharts) as well. Because they’re separate elements, you can easily change the structure of a PivotTable, too, and your changes won’t mess up the data in the underlying table. Using interactive web components, you can also make PivotTables available to other people via a web browser.

→ For details on how to use PivotTables in web pages, see “Using Office Programs to Create and Edit Web Pages,” p. 59.

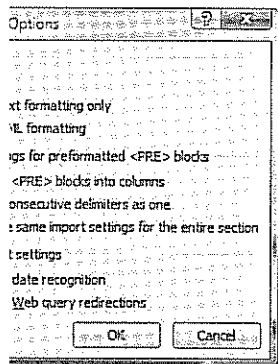
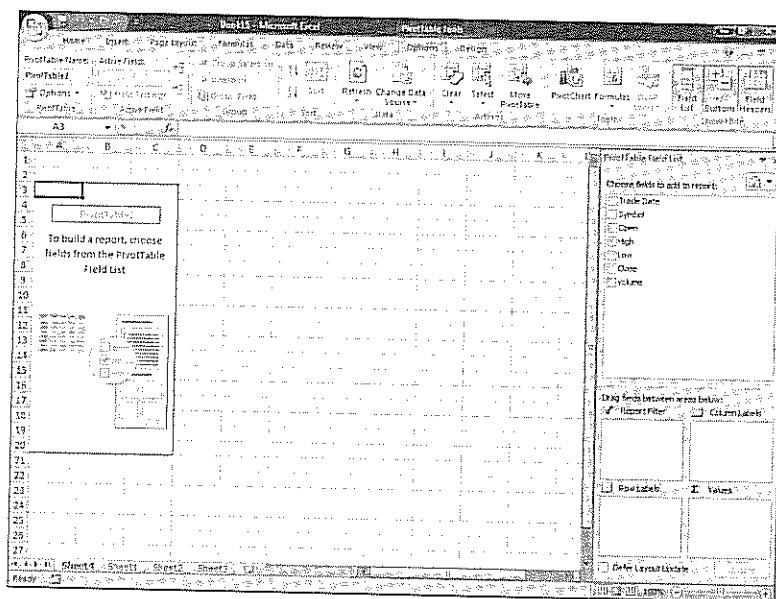


Figure 21.27 shows the PivotTable Task pane that appears when you elect to create a PivotTable from a table. Excel 2007 makes it simple to create a PivotTable; you just choose data fields from the PivotTable field list in the Task pane to add data to your PivotTable report.

The PivotTable Field List includes every field in your table. Use row fields and column fields to define how you want Excel to group your table. Data items define which fields contain the information you want to summarize. Page fields let you further refine your view by displaying a separate PivotTable for each item in a group, as though the table were on its own virtual page. You can use multiple row fields, column fields, or both, and you can specify which summary action you want Excel to perform on data items—the sum, average, or count of all related values, for instance.

Figure 21.27
You'll use the
PivotTable Task pane
to manage your
PivotTable.



What can you do with a PivotTable? The number of uses is limited only by your imagination. Despite their dramatically different structures, for example, each of the following four PivotTables started with the same table of information about publicly traded stocks. In its raw form, with its grand total of 106,224 separate data points, the table is a prescription for information overload. Each of the 6,639 rows contains 16 data fields for an individual publicly traded company, including its name, ticker symbol, and industry category, the exchange on which it trades, its high and low stock price for the past year, and financial measurements such as net profit margin and return on equity.

Figure 21.28 shows a simple PivotTable that lets you see at a glance how many companies are in each industry category, along with the average increase or decrease in stock price from companies in that category over the past year. This PivotTable consists of a single row field and two data items.

elect to create a PivotTable; you just choose a field to add to your PivotTable

fields and column labels to define which fields control the view of the table. You can refine your view by adding or removing fields, and you can specify the calculation, such as the sum, average, or

Figure 21.28
With no column fields and only one row field, this PivotTable quickly counts the number of companies in each category and calculates the average price change for the year.

	A	B	C
	Industry Name	Number of Companies	% Price Change Last Year (Avg)
1			
2	Closed-End Fund-Debt	360	(1.00)
3	Savings & Loans	237	18.94
4	Business Software and Services	164	(38.23)
5	Internet Software and Services	133	(22.66)
6	Business Services	128	(21.35)
7	Biotechnology	120	(45.73)
8	Scientific & Technical Instruments	119	(24.45)
9	Regional - Northeast Banks	116	14.96
10	Drug Manufacturers - Other	111	(33.29)
11	Independent Oil & Gas	106	(7.48)
12	Medical Appliances & Equipment	103	(18.11)
13	Regional - Mid-Atlantic Banks	99	19.95
14	Communications Equipment	94	(31.54)
15	Medical Instruments & Supplies	88	(24.90)
16	Application Software	85	(27.05)
17	Property & Casualty Insurance	75	(13.30)
18	Restaurants	71	(19.31)
19	Closed-End Fund-Foreign	71	(11.76)

In Figure 21.29, more detail is added, displaying individual statistics for each company and grouping the detail rows in alphabetical order by industry name. For this PivotTable, the data is arranged in report format, similar to the banded database reports Access and other database management programs produce. Note that this PivotTable includes four data items instead of two, and a slew of Excel formatting options are used to make the report more readable—changing fonts and font sizes, aligning type and adding background shading, and standardizing the number of decimal points in each column.

Figure 21.29
To hide gridlines and group-related items in bands such as these, choose a report format instead of the default table layout.

	Industry Name	Company Name	Avg of Net Profit Margin	Avg of Service Charge Last	52-Week Low	52-Week High
127		Petroleum Helicopters, Inc.		5.8	23.9	31
128						
129						
130		Akzo, Inc.	3.0	-25.1	23.2	40.1
131		Akzo, Inc. 2nd	1.9	-41	17.6	39.1
132		ALUMINUM CHLORIDE	9.8	-17.5	9.4	22.0
133		Century Aluminum Company	(2.6)	-50.0	5.5	17.5
134		Commonwealth Industries, Inc.	0.9	-26.2	-4.2	8.1
135		Pechiney	0.8	-46.6	11.4	27.8
136						
137						
138		Abercrombie & Fitch Co.	12.2	0.7	15.0	31.9
139		American Eagle Outfitters, Inc.	6.1	-41.1	9.8	27.2
140		Ann Taylor Stores Corporation	5.8	-28.1	17.1	33.2
141		Arden Stores, Inc.	7.4	-45.3	9.6	24.0
142		Bag-Tag Holdings, Inc.	3.4	-35.0	1.8	5.6
143		Buckle, Inc.	8.0	-26.4	15.5	25.5
144		Burlington Coat Factory Warehouse Corp.	2.7	-11.7	15.4	23.5

→ For details on how to create similar reports from an Access database, see "Building Forms and Reports," p. 852.

To slice the data even more finely and add an extra analytical dimension, you can drag more items from the PivotTable Task pane to the row and column fields. Each row in the PivotTable is grouped using unique values in two categories, and there are two column headings as well, one for each unique value in the Split in Last Year column field. (To make the PivotTable easier to read, the column headings were renamed from Yes and No to Split

and No Split.) At the intersection of each row and column in the PivotTable, Excel counts the number of companies and calculates the average income per employee for all rows that match the row and column fields.

The resulting PivotTable, shown in Figure 21.30, is a concise and crystal-clear cross-tabulation, giving you a side-by-side analysis of the number of stocks that split in the past year versus those that didn't, broken down by industry category and exchange.

Figure 21.30
Add a column field to quickly compare related data points. Notice that the worksheet pane is frozen to keep headings visible when scrolling, just as with an ordinary worksheet.

Industry Name	Exchange	Split in Last Year		No Split	
		Number of Companies	Avg Income Per Employee	Number of Companies	Avg Income Per Employee
140 Drugs Wholesale	AMEX	0	\$ 0	2	\$ (174,500)
141	NASDAQ	1	\$ 7,000	2	\$ 54,000
142	NYSE	1	\$ 22,000	2	\$ 17,000
143 Education & Training Services	NASDAQ	1	\$ 21,000	3	\$ 16,500
144	NYSE	1	\$ 12,000	2	\$ 4,000
145 Electric Utilities	AMEX	0	\$ 0	3	\$ 35,000
146	NASDAQ	0	\$ 0	2	\$ 16,000
147	NYSE	3	\$ 53,333	14	\$ 31,619
148 Electronic Equipment	AMEX	0	\$ 0	3	\$ 21,000
149	NASDAQ	1	\$ 34,000	5	\$ (400)
150	NYSE	0	\$ 0	1	\$ 12,000
151 Electronics Stores	NASDAQ	0	\$ 0	3	\$ 0
152	NYSE	1	\$ 6,000	3	\$ 1,667
153 Electronics Wholesale	NASDAQ	0	\$ 0	5	\$ 6,200
154	NYSE	0	\$ 0	2	\$ 2,000
155 Entertainment - Diversified	AMEX	0	\$ 0	2	\$ 6,000
156	NYSE	2	\$ 101,000	4	\$ 24,500
157 Farm & Construction Machinery	AMEX	1	\$ 12,000	0	\$ 0
158	NASDAQ	0	\$ 0	2	\$ 12,500

There are literally hundreds of options in even a modestly complex PivotTable, but a PivotTable doesn't have to be large or complex to be effective. The PivotTable in Figure 21.31, for example, neatly summarizes more than 100,000 data points in just a few rows and columns.

Figure 21.31
Notice the grand totals under the rows in this PivotTable. Use the page field in the top-left corner to filter the entire table.

Exchange	Dow Jones Membership	Net Profit Margin (Avg)	Price Change Last Yr (Avg)
AMEX	(blank)	0.85	-10.85
NASDAQ	DI Industrials	21.25	-25.45
	DI Transports	-0.46	-12.93
	(blank)	-3.35	-17.26
NYSE	DI Industrials	8.43	-21.23
	DI Transports	2.05	-27.34
	DI Utilities	-4.83	-34.26
	(blank)	6.29	-16.25
Grand Total		0.56	-16.32

To produce this example, we used two column fields, two row fields, and one page field—a drop-down table that lets us filter the records in the entire table. Choosing (All) from the page field shows a summary of all data in the table. By selecting a different entry from the drop-down table, you can show the same breakdown for each industry name. Select one category at a time to flip through a series of otherwise identical PivotTables that focus on each category.

The layout Excel produced automatically included totals for each row and column; we kept only the grand total at the bottom of the PivotTable. We had to modify other default settings as well, including changing the default formula to calculate the average of our data items. To make the headings and totals easier to read, we did some rewording, and then changed fonts and alignment, added shading, and wrapped text.

PivotTable, Excel counts employee for all rows that

crystal-clear cross-tabulation split in the past year version.

Avg Income Per Employee
\$ (174,500)
\$ 54,000
\$ 17,000
\$ 16,500
\$ 4,000
\$ 55,000
\$ 10,000
\$ 31,515
\$ 21,000
\$ (400)
\$ 12,000
\$ 0
\$ 4,567
\$ 6,200
\$ 2,000
\$ 6,000
\$ 24,500
\$ 0
\$ 10,500

PivotTable, but a PivotTable in Figure 21.32 in just a few rows and

and one page field—a choosing (All) from the different entry from the category name. Select one categories that focus on each

view and column; we kept the other default settings—average of our data and wording, and then

WHEN SHOULD YOU USE A PIVOTTABLE?


PivotTables have several advantages over other worksheet models. Using the PivotTable button, it's easy to create a PivotTable that summarizes all or part of a table in dozens of different ways. Trying to accomplish the same task by entering formulas manually would take days. Also, because PivotTables and PivotCharts do not change your existing data or its arrangement on the worksheet, you can freely experiment with different PivotTable layouts. Use the Undo button (on your Quick Access Toolbar or use the Ctrl+Z shortcut key) to roll back any changes you make in a PivotTable layout. If you want to start over, you can delete the PivotTable page and run the wizard again.

PivotTables are the correct choice when all your data is in a table or in an external database that you can query from Excel. PivotTables are not appropriate for structured worksheet models that include data-entry cells, subtotals, and summary rows. A PivotTable won't do much good on an annual budget worksheet, for example, because it already includes rows, columns, and subtotals. On the other hand, if you enter the raw data in a table (or import it from an external database), with each row containing a month, department, budget category, and amount, you can easily re-create that same layout in PivotTable form—and you'll have many more analytical options available to you later.

→ For more details on how to use Microsoft Query to pull data from an external database, see "Connecting a Worksheet to an External Databases," p. 645.

CREATING A PIVOTTABLE

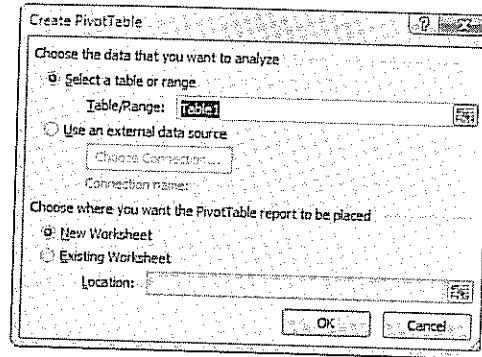
To create a PivotTable from an existing table, you'll use the Insert ribbon's PivotTable button. Excel will prompt you for basic details about the PivotTable you want to create, including the location of the data source and where you want the PivotTable to appear. After you finish specifying the PivotTable, you'll be able to lay out your data directly on the worksheet.

 You don't need to use a PivotTable Wizard in Excel 2007 as you did in Excel 2003. The PivotTable screen is better formatted to help you build and edit your PivotTables and the simple Insert PivotTable dialog box is the only place you need to specify your initial PivotTable data source.

To build a new PivotTable, open the workbook that contains the table on which you plan to base the PivotTable. Then follow these steps:

1. Click anywhere in your table. To build a PivotTable from a subset of the data in your table, select the range that contains the data.
2. Click your Insert ribbon's PivotTable button. The Create PivotTable dialog box appears, as shown in Figure 21.32.

Figure 21.32
Select a range or use
an external data
source for your
PivotTable



Using External Databases with PivotTables

In a corporate setting, it's often useful to base a PivotTable on the result of a query to an external database. If you choose the External Data Source option, Excel starts the Query Wizard and prompts you for details about the format and location of the database. Excel then uses this query as the source for the PivotTable. Each time you refresh the data in the PivotTable, Excel runs the saved query and updates the PivotTable with the most recent information.

Excel also offers the option to build PivotTables from special data structures called On-Line Analytical Processing (OLAP) databases. Instead of rows and columns, these files organize data into dimensions and levels. Instead of forcing Excel to chug through massive amounts of data, the server does the summarizing first and sends the summary values directly to your report.

When you connect to an OLAP database, Excel lets you save your data in local files called *OLAP cubes* and use them as the source for a PivotTable. There are some substantial differences in the way PivotTables based on OLAP data work compared with those based on Excel tables or non-OLAP databases.

To learn more about building PivotTables from an OLAP database, pick up a copy of *Special Edition Using Microsoft Office Excel 2007*, also published by Que.

3. Specify the range in which your data is located. The default selection is your current table, or any range you selected before requesting the PivotTable. Adjust the selection, if necessary.
4. Specify where you want to place the PivotTable. Choose the default option, New Worksheet.

CAUTION

The Create PivotTable dialog box offers the option to place a PivotTable on an existing worksheet. In general, you should always choose to place a PivotTable on its own sheet. Adding a PivotTable to a sheet that contains data exposes you to the risk that changes you make to the table design will affect your PivotTable, or vice versa.

5. Click OK to close the Create PivotTable dialog box and create a blank PivotTable page. Excel jumps to the new worksheet you just created and displays the PivotTable Field List in the Task pane to the right of your screen.

6. Drag field buttons from the Choose Fields to Add to Report box and drop them into the appropriate regions below to one of the four regions: Report Filter, Column Labels, Row Labels, or summed Values.

TIP FROM

EQ & Woody

Click the drop-down arrow in the top-right corner of the Field List Task pane to change the Task pane's layout if the default setting makes the targets too far away to drop the fields. If you're working with massive data, you might want to click the Layout Update option so Excel doesn't rearrange its tables as you build your PivotTable. Click Update whenever you want to see the iterative result of your work.

Don't be surprised if the PivotTable doesn't display properly at first. In particular, summary fields in the Values area default to the SUM function. If you want to use COUNT, AVERAGE, or another summary function instead, see the next section.

EDITING AND UPDATING A PIVOTTABLE

Now that you've seen the PivotTable build process, it will be helpful to review the actions you can take with PivotTables as you build and edit them. The following list reviews the drag-and-drop operations you'll perform as you drag data fields from the field list to the four PivotTable destinations:

1. Drag a data field to the Report Filter area at the top of the PivotTable report you build.
2. Drag a data field to the Column Label area to make that field your column labels that run across your PivotTable.
3. Drag a data field to the Row Label area to make that field your row labels that run down the left side of your PivotTable report.
4. Drag a data field to the summed Values area to make that field the data that composes the body of your PivotTable report.

The location where you drag a data field, therefore, determines where on the PivotTable report that field will make itself be known. After items begin to appear on your PivotTable report, tablelike drop-down arrows enable you to select and sort an item as needed. You can use the same data field in more than one location.

You can add fields to the PivotTable report by selecting the check box next to each field name. Excel adds nonnumeric fields to the Row Labels area, numeric fields to the Values area, and OLAP date and time hierarchies (if you use that) to the Column Labels area. At any point you can right-click a field name and select an Add To command to add that field to an area without using a drag-and-drop operation.

After you create a PivotTable, it's easy to rearrange fields and data items. Drag fields from one place to another to change the display of data—from a row field to a column field, for example, if you want to see values side by side rather than one above the other. Right-click to display shortcut menus that let you adjust formatting and other options for each field.

If the Field Table isn't visible, right-click the PivotTable layout area and choose View Field Table from the bottom of the shortcut menu. To make changes to the PivotTable report, use any or all of these techniques:

- To add a new field to the layout, drag a field button from the PivotTable field list and drop it on the layout. If you're replacing an existing field, remove the old field first to reduce unnecessary calculations. When you drop a new field in the row or column area, Excel adds it as part of the hierarchy of fields that are already there and automatically groups items in the order in which they appear. Be careful to arrange these fields in the proper order. For example, if you have a table of product categories, each of which contains multiple products, place the category field to the left of the product name field, or the results will be nonsense.

TIP FROM

EQ: Woody

If your table includes two fields that have an absolute one-to-one correspondence, such as part numbers and part names, you can add them to the row area in either order and your table will appear correctly.

- To remove a field from any part of the PivotTable layout, drag the field button off the layout; when the pointer icon changes to include an X, release the mouse button. You can also click a dropped field's down arrow that appears to the right of a field name and select Remove Field.
- To change the order of fields in rows, columns, or the data area, drag the field button and drop it in the correct location on the layout. Make sure you're pointing to the field button and not its label; you'll know you've aimed correctly when the mouse button turns to a four-headed pointer. Drag to another location and watch the mouse pointer and thick black lines for feedback on the correct "drop" location.

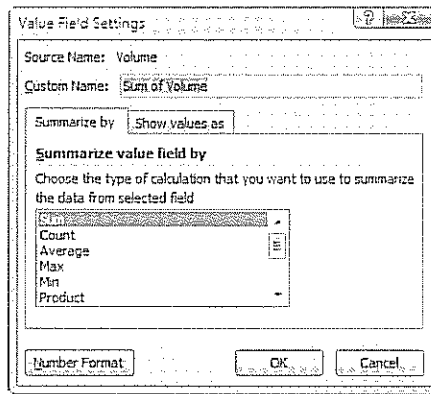
TIP FROM

EQ: Woody

Using the mouse to rearrange the order of data items on a PivotTable can be frustrating. It's usually easier to right-click the field button you want to move, and then choose any of the options on the Order menu. Typically, you can move the item left or right one position, or move it to the beginning or end of the table.

- To change the summary function used in the data area (from SUM to COUNT or AVERAGE, for example), right-click the field button in the PivotTable and choose Value Field Settings from the shortcut menu. That action opens the Value Field Settings dialog box, shown in Figure 21.33. Select a function from the Summarize Value Field By; if you want to change the name from its default, do so in the Custom Name box, and then click OK to save the change.

Figure 21.33
Change the summary
you wish to see.



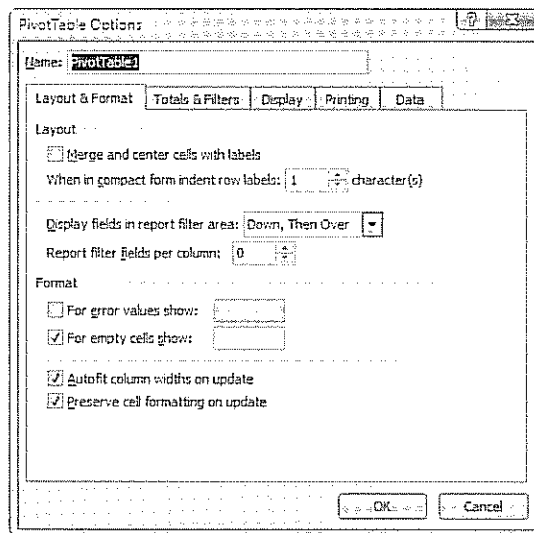
When you drag and drop buttons to arrange fields on a PivotTable page, Excel makes all kinds of decisions on your behalf. If these defaults aren't correct, the following sections will help you change them.

CHANGING A PIVOTTABLE'S APPEARANCE

The default sort order for rows and columns is usually alphanumeric. You can change the order of individual items by dragging them up or down (in the case of rows) or left or right (for columns). In other cases, you might want to adjust the default sort order. For example, if your PivotTable counts the number of items in each category, you might want to see categories with the highest number of items at the top of the table.

Right-click any PivotTable report item and select PivotTable Options. Excel opens the PivotTable Options dialog box shown in Figure 21.34.

Figure 21.34
The PivotTable
Options dialog box
enables you to
change multiple
aspects of your
PivotTable data.



Click the arrow to the right of any PivotTable button for the row or column field and choose More Sort Options to display the Sort dialog box that enables you to modify the order of your data. Click the Sort dialog box's More Options button to set additional sort options, such as whether you want to sort every time a report is updated (if your data is massive, you'll want to uncheck this option).

Click the down arrow to the right of any row or column label and select Label Filters or Value Filters to limit data that appears in the report. The options you see, such as Greater Than and Between and Top 10 work like table AutoFilter options (because that's what they actually are). This is a good way to create a "top 10" table, for example, showing only the categories that have the most items.

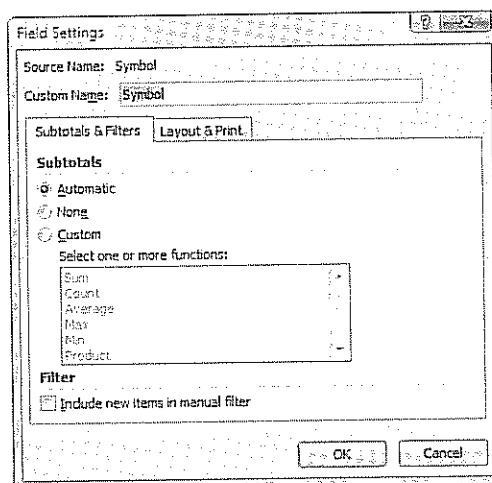
→ AutoFilter can save a tremendous amount of time, if you know how to use it properly; see "Using Filters to Find Sets of Data," p. 634.

ADDING AND REMOVING SUBTOTALS

You can add subtotals to rows, columns, or both in a table. In some cases, Excel adds them automatically, even if they're not appropriate. Subtotals can add a useful way to see the impact of groupings in your PivotTable, or they can add clutter between rows and columns. Depending on the design of your PivotTable and what Excel did automatically, you might need to add or remove these subtotals. In some cases, you can remove subtotals with the right-click shortcut menu. Right-click any of the subtotals and choose Hide. To add subtotals, you need to use the dialog boxes. To work with subtotals, follow these steps:

1. Right-click any row or column heading that contains the subtotal, and choose Field Settings from the shortcut menu. Excel displays the Field Settings dialog box, as shown in Figure 21.35.

Figure 21.35
Use the Subtotals options to add, edit, or hide subtotals for a row or column.



2. In the Subtotals section, choose Automatic to let Excel create subtotals for all items. Choose Custom and click a summary function to add one or more specific type of subtotals, such as Count and Average. Click None to remove all subtotals.
3. Click OK to exit the dialog box and make the changes you specified.

REMOVING BLANK CELLS AND ERROR MESSAGES

Because PivotTables automatically summarize all data, it's common to see blank cells and error messages in the data area. #DIV/0 errors, for example, are especially common when calculating averages because in a long table, it's almost certain that some items will have no matches in a particular row-and-column intersection. For example, if you're calculating average sales with regions in the column area and product categories in the row area, some regions will have no sales for a particular category. These aren't really errors; instead, you want the table to display a label such as NA, for Not Applicable.

Careful attention to blanks and error messages can make your PivotTable easier to read and make it look more professional. Here's how to adjust the appearance of blank cells and errors:

1. Right-click any part of the PivotTable and choose Properties from the contextual menu.
2. Select the For Error Values, Show check box. Click in the box to the right and fill in the information you want to display instead of the error message, such as NA.
3. Select the For Empty Cells Show check box. If the field contains numeric data, enter 0 here; for a text field, enter the value you want Excel to display (NA, for instance) instead of leaving the cell blank.
4. Click OK to save your changes.

CHANGING OR REFRESHING PIVOTTABLE DATA

When you change the layout of a PivotTable, Excel automatically recalculates the resulting display of data. If you add or edit data in the underlying table, however, your changes do not appear immediately in the associated PivotTable. For PivotTable reports based on Excel tables, you must manually refresh the data in the PivotTable whenever you add, remove, or edit data. To be certain that the PivotTable reflects all recent changes, click the Update button on the PivotTable Task pane.

TROUBLESHOOTING

SORTING OUT SORTING PROBLEMS

I tried to run a multicolumn sort on my table, but the result came out scrambled and my column labels disappeared.

The most likely cause of this problem is that Excel couldn't identify the header row in your table. If you use the Sort dialog box, you can select a check box that tells Excel that your table has a header row. Other common sorting problems are caused by blank rows or columns in the table, in which case Excel doesn't sort the records below the blank row or to the right of the blank column. To work around this problem, select your table, minus the header rows, and choose the No Header Row option. You'll need to make a mental note of which column number corresponds with each column.

